

# Initializing adaptive importance sampling with Markov chains

Frederik Beaujean\*, Allen Caldwell

*Max Planck Institute for Physics*

---

## Abstract

Adaptive importance sampling is a powerful tool to sample from complicated target densities, but its success depends sensitively on the initial proposal density. An algorithm is presented to automatically perform the initialization using Markov chains and hierarchical clustering. The performance is checked on challenging multimodal examples in up to 20 dimensions and compared to results from nested sampling. Our approach yields a proposal that leads to rapid convergence and accurate estimation of overall normalization and marginal distributions.

**Keywords:** adaptive importance sampling, population Monte Carlo, Markov chain, hierarchical clustering, multimodal

---

## 1. Introduction

The fundamental problem we wish to address is to sample from or integrate over a complicated density  $P(\theta)$ , the *target density*, in a moderately high-dimensional parameter space. Our main application is Bayesian statistics, where  $P(\theta)$  is identified with the unnormalized posterior distribution. The samples are useful for parameter inference, and for model comparison it is necessary to compute the posterior normalization—the evidence or marginal likelihood,  $Z$ —given by the integral over the product of likelihood  $L(\theta)$  and prior  $P_0(\theta)$  as

$$Z = \int d\theta P(\theta) = \int d\theta L(\theta)P_0(\theta). \quad (1)$$

A plethora of sampling algorithms exists in the literature; a comprehensive, though somewhat dated, overview is presented in [1]. In contrast to algorithms tailored to very specific targets, we rather want to make progress toward the “holy grail” of sampling: an efficient, accurate, and parallelizable algorithm that copes with any (continuous) target in any dimension. Ideally, such an algorithm yields samples and the integral in one run.

Current analyses at the frontier of cosmology and elementary particle physics often involve extensions of the accepted standard models with a large number of parameters, but only loose constraints from the available data; see, for example, [2, 3]. The methods presented here were developed in the course of a global analysis of rare B-meson decays described in detail in [4, 5]. Posterior densities occurring in that analysis exhibit many of the typical features that require a sophisticated sampling procedure as there are degeneracies due to continuous symmetries and multiple isolated modes arising from discrete symmetries in  $d = 18 - 31$  dimensions. The standard approach based on local random-walk Markov chains (MCMC) is notorious for failing to produce reliable estimates under these circumstances as chains mix very slowly or not at all.

In challenging problems with large data sets or difficult-to-obtain model predictions, the sampling is further complicated by the relatively long time required to evaluate the target density (roughly 1 s in the motivating example [4]). Apart from more efficient algorithms, the most straightforward option to reduce the wallclock time needed for the computation is to use parallel computing facilities. The easiest implementation is done on a standard system with

---

\*Corresponding author

Email addresses: beaujean@mpp.mpg.de (Frederik Beaujean), caldwell@mpp.mpg.de (Allen Caldwell)

multiple cores, but it is desirable to use computing clusters or graphical processing units offering hundreds or even thousands of cores.

Adaptive importance sampling, or population Monte Carlo (PMC) [6, 7], is an evolution of classic importance sampling [1, 8] combining most of the desirable features outlined above. The basic idea is to use a mixture density as a proposal function  $q(\theta)$ , and to iteratively update  $q(\theta)$  to match the target density as closely as possible. The individual components of  $q(\theta)$  are conveniently chosen as multivariate Gaussian or Student’s  $t$  distributions; by adjusting their location and covariance one can easily accommodate multiple modes and degeneracies. Each proposal update requires a step where a large number of i.i.d. samples  $\{\theta^i : i = 1 \dots N\}$  are drawn from  $q(\theta)$ , allowing trivial parallelization of the potentially costly evaluation of the importance weights  $w_i = P(\theta^i)/q(\theta^i)$ .

The PMC update algorithm is based on expectation-maximization (EM) [9] and seeks to reduce the Kullback-Leibler divergence [10] between target and proposal,  $\text{KL}(P\|q)$ . Since  $\text{KL}(P\|q)$  usually has multiple minima, and EM tends toward a local minimum, the initial guess for the proposal is of utmost importance for the success of the algorithm. With a poor proposal, PMC fails as proposal updates lead to a consecutively poorer approximation of the target. In that case, typically more and more of the proposal components are assigned vanishing weight until only one remains. Since a single component is insufficient in all but the most trivial cases, the PMC results are useless. The authors of PMC and its reference implementation [11] in the C language commented on the issue of the initial guess, but provided only basic advice that is useful just in fairly simple unimodal problems. In their first applications of PMC to physics analyses [12, 13], it was sufficient to scatter mixture components around the center of the parameter range or around a previously computed maximum using the Fisher matrix and “educated guesses” [13]. However, this approach did not give satisfactory results for the analysis presented in [4].

Clearly, a more robust approach to initialization is preferable. Previous attempts at such an initialization were suggested in the context of econometrics [14] and population genetics [15]. In the basic algorithm of [14], the authors propose to start with a single component given by the mode and the inverse Hessian at the mode. In each update, one new component is added, which is constructed from the highest importance weights. The algorithm terminates when the standard deviation of the weights divided by their mean is sufficiently small. For multimodal problems, they suggest a tempering approach to first adapt the proposal to a simplified target and demonstrate successful discovery of 20 well separated Gaussians, albeit only in 2D. For nonelliptical target densities in higher dimensions where many dozens of components are needed, their approach would presumably require an excessive number of updates. In [15], a large sample from the uniform or prior distribution with a logistic rescaling is used to learn the features of the target. The authors propose to run Gaussian mixture clustering with the integrated likelihood criterion fixing the optimal number of components of the initial proposal for PMC. By cleverly combining the samples of *all* PMC update steps, and not only the most recent one as in our approach, they report a significant Monte Carlo variance reduction. Nested sampling [16] is an alternative technique to simultaneously compute weighted samples and the normalization of a complicated target density. The basic idea is to evolve a collection of sample points such that in each iteration the point with the smallest likelihood is replaced by a new point with larger likelihood drawn from the prior. Multinest [17] is the most widely used implementation of nested sampling; it uses sets of ellipsoids to map out the target’s regions of interest.

Our approach is similar to the efforts of [14, 15] in that we seek to create a good initial proposal for PMC with a minimum of manual intervention. The initialization proceeds in two phases. In the learning phase, multiple local random-walk Markov chains are run in parallel to discover and explore the regions where  $P$  is large. In the next phase, we use the chains to extract the local features of the target by partitioning up chains into short patches of around 100 iterations and thus turn one of the weaknesses of the random walk—the slow diffusion-like exploration—into a virtue. Each patch defines a mixture component through its sample mean and covariance. To reduce the number of components to a tractable number, we employ hierarchical clustering [18]. The proposed initialization differs from previous attempts in its usage of Markov chains and is designed specifically for complicated targets in fairly high dimensions  $d \lesssim 40$  with complicated shapes such as degeneracies, multiple modes, and other nonelliptical structures.

After a detailed description of the algorithm in Section 2, and a brief summary of its parameters in Section 3, we illustrate the algorithm with several examples in various dimensions in Section 4 and compare its performance to that of Multinest version 2.18. We do not compare to alternative PMC initializations since no implementations were available to us when this work was carried out. Ideas on future directions and concluding remarks are presented in Sections 5 and 6.




Figure 1: Overview of the algorithm described in Section 2.1. The different  $q$  refer to the mixture densities used at different stages of the algorithm, while the  $K$  refer to the number of components.

## 2. Algorithm

### 2.1. Overview

Our focus is on creating a good initial proposal for PMC with a minimum of manual intervention. In general, it is necessary to

1. explore the target  $P$  through evaluations at a number of points in parameter space; and to
2. extract and combine the knowledge from the exploration into a mixture density that approximates  $P$ .

In practice, our suggestion is to *combine the best of MCMC and PMC* in three steps (cf. Fig. 1):

1. We run  $k$  Markov chains in parallel, each performing a local random walk with an adaptive proposal on the target  $P(\theta)$ .
2. Then we extract the support of the target density by splitting each chain into many small patches. Sample mean and covariance of each patch define one multivariate density, the collection of  $K_{\text{patches}}$  patches yields a mixture density  $q_{\text{patches}}(\theta)$ .
3. Typically, there are more patches than actually needed; hierarchical clustering produces a mixture with far fewer components  $K \ll K_{\text{patches}}$  but essentially the same knowledge of the target density by removing redundant information. The output of hierarchical clustering,  $q_{\text{HC}}(\theta)$ , is used with minor modifications as the initial proposal for PMC. We then run the standard PMC updates implemented in [11] until convergence and extract samples  $\theta^i$  with importance weights  $w_i$  using the proposal  $q_{\text{final}}(\theta)$ .

The combination of MCMC and PMC is one leap forward towards a black-box Monte Carlo sampler that learns the relevant features of the target density automatically. In order to make optimal use of a parallel computing infrastructure, the MCMC prerun is to be kept at a minimum length, and preferably most evaluations of the target density are performed during the PMC phase, when massive parallelization is available.

### 2.2. Detailed description

#### 2.2.1. MCMC prerun

For concreteness, we chose to implement the adaptive Metropolis algorithm suggested in [19] for the examples discussed below. It uses a multivariate Gaussian proposal that is centered on the current point and updated based on the covariance of previous iterations with the cooling scheme described in [12]. Note that the resulting chain strictly speaking is not a Markov chain because the proposal is continuously adapted after every batch of  $N_{\text{update}}$  iterations, but for simplicity, we continue to use the term “Markov”. In our algorithm, we only rely on the fact that the samples are generated by a local random walk and that their asymptotic distribution is  $P(\theta)$ .

Assuming no knowledge of the target density other than that is zero outside of a given hyperrectangle in  $\mathbb{R}^d$ , we draw the initial positions of the chains from the uniform distribution. If the target is a posterior density and the priors are of a simple form, we can draw the starting points directly from the prior. Similarly, the initial covariance matrix is proportional to

$$\Sigma^0 = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_d^2), \quad (2)$$

where  $\sigma_i^2$  is the prior variance of the  $i^{\text{th}}$  parameter. We then rescale  $\Sigma^0$  by  $2.38^2/d$  [20] to increase the efficiency of the initial proposal; in subsequent updates, we update the scale factor to achieve an acceptance rate between (15 – 35) %.

In most problems, the prior is significantly more diffuse than the posterior, hence our choice of seeding the chains automatically generates overdispersion. The main reason why overdispersion is desirable to us is that in the face of potentially several maxima, with little analytical knowledge of a posterior that is often available only in the form of computer code, it is imperative to explore the full parameter space and to find *all* regions of significant probability. These regions are not limited to local maxima, but include degenerate regions as well. Therefore, the number of chains,  $k$ , should be chosen significantly larger than the number of expected maxima. If the location of the maxima is known, an equal number of chains can be started in each maximum for higher efficiency. Since in many realistic problems the purpose of the sampling is to discover the maxima, in the examples below we choose *not* to use the available analytical knowledge on the location of the modes in order to highlight potential pitfalls.

We select a fixed number of iterations,  $N_{\text{MCMC}}$ , to terminate the sampling without regard for chain mixing.  $N_{\text{MCMC}}$  ought to be chosen rather small in the trade-off between accuracy and computing time. Even in the most complicated settings that we treated [5],  $N_{\text{MCMC}} \lesssim 40\,000$  revealed enough information about  $P(\theta)$ .

Given the prerun of  $k$  chains, we extract the local information by exploiting the slow, diffusion-like exploration of the chain. To this end, we choose a *patch length*  $L$ , and partition the history of each chain, with the exception of the burn-in, into patches of length  $L$ . For the  $i^{\text{th}}$  patch, we compute the sample mean  $\mu_i$  and sample covariance  $\Sigma_i$ , and form a multivariate Gaussian density. Patches in which no move is accepted are discarded, and those for which the numerical Cholesky decomposition fails are used with off-diagonal elements of the covariance matrix set to zero. The patch length ought to be chosen in such a way that small-scale features of the posterior can be explored during  $L$  iterations. A good value of  $L$  slightly increases with  $d$  and possible degeneracies. On the other hand,  $L$  must not be too small, else the chain cannot move enough. Combining patches from all  $k$  chains, we obtain a Gaussian mixture density of  $K_{\text{patches}}$  components

$$q_{\text{patches}}(\theta) \equiv \sum_{i=1}^{K_{\text{patches}}} \alpha_i \mathcal{N}(\theta | \mu_i, \Sigma_i). \quad (3)$$

and assign equal weight  $\alpha_i = 1/K_{\text{patches}}$  to each component. Note that we do *not* take into account the value of the posterior in each patch; rather, we will ultimately rely on PMC to find the proper component weights.

### 2.2.2. Hierarchical clustering

The information about the target from the Markov chains is contained in a mixture density with a large number of components  $K_{\text{patches}}$ , where typical values of  $K_{\text{patches}}$  can reach several thousands. For computational efficiency, it is important to reduce the complexity of the mixture to keep the number of samples needed in each PMC update step low. At the same time, we wish to preserve as much information as possible. The goal is to compress the  $K_{\text{patches}}$  components into a mixture with only  $K \ll K_{\text{patches}}$  components by removing redundant information that, for example, comes from multiple chains that mix or from a single chains repeatedly visiting a region. *Hierarchical clustering* [18] is our weapon of choice. It achieves the compression by finding a  $K$ -component Gaussian mixture

$$q_{\text{HC}}(\theta) \equiv \sum_{i=1}^K \alpha_i \mathcal{N}(\theta | \mu_i, \Sigma_i). \quad (4)$$

that minimizes a distance measure based on the Kullback-Leibler divergence [10].

#### Initialization

Hierarchical clustering, being an expectation-maximization variant, converges only on a local minimum of the distance measure. Given a large number of input components, there exist numerous local minima, hence it is crucial to supply good initial guesses for the output components, such that the initial solution is already very close to a *good* final solution. We then note rapid convergence after  $\mathcal{O}(10)$  steps. There are two important questions to address:

1. Where to put the initial output components?
2. How many output components,  $K$ , are needed?

At present, we assume a fixed value of  $K$ . In [18], it is vaguely recommend to use “standard methods for model selection” to determine  $K$ . We can only speculate that they refer to the Bayesian information criterion [21] or the

Akaike information criterion [22]. Another approach would be to add one component at a time until  $K$  is “large enough”. It then remains to specify a quantitative stopping criterion. In [14], an attempt at such a criterion is presented, but it appears somewhat inefficient when a large number of components is needed. As a rule of thumb, we recommend  $K$  should be at least as large as  $d$ .

But to answer the first question, we have a good idea *where* to place the components. The key is to group the chains, and to have a fixed number of components per group from *long patches*; i.e., parts of chains of length significantly exceeding  $L$ . To begin with, it is necessary to determine which chains have mixed in the prerun. Two or more chains whose common Gelman-Rubin  $R$  values [23] are less than a given constant,  $R_c$ , for all parameters, form a *group* of chains. Most importantly, this ensures that a similar and sufficient number of components is placed in every mode of the target density, regardless of how many chains visited that mode. We ignore the burn-in samples of each chain as described in the previous section.

Let us assume we want  $K_g$  components from a group of  $k_g$  chains. If  $K_g \geq k_g$ , we find the minimal lexicographic *integer partition* of  $K_g$  into exactly  $k_g$  parts. Hence, the partition, represented as a  $k_g$ -dimensional vector of integers  $\mathbf{n}$ , is given by

$$\mathbf{n} = \left( \left\lceil \frac{K_g}{k_g} \right\rceil, \dots, \left\lceil \frac{K_g}{k_g} \right\rceil, \left\lfloor \frac{K_g}{k_g} \right\rfloor, \dots, \left\lfloor \frac{K_g}{k_g} \right\rfloor \right), \quad (5)$$

where we use the ceiling ( $\lceil \cdot \rceil$ ) and floor ( $\lfloor \cdot \rfloor$ ) operations. The first  $(K_g \bmod k_g)$  parts are one larger than the remaining parts. For example, with  $K_g = 6$  and  $k_g = 4$ , the partition is  $(2, 2, 1, 1)$ . In the rare case  $K_g < k_g$ , the integer partitioning cannot be performed as above. Instead, we combine all individual chains into one long chain, and set  $k_g = 1$ .

Finally, the  $i^{\text{th}}$  chain is partitioned into  $n_i$  long patches, and the sample mean and covariance of each patch define one multivariate Gaussian. The long patches, say there are two or three per chain, represent expectation values over many iterations. Small-scale features are averaged out, while the center of gravity is preserved. Thus the initial output components from one group are very similar, and the hierarchical clustering shifts and shrinks them to fit. It is possible that some of the components are assigned zero weight, but due to the initial similarity, very few, and usually zero, components “die” during the hierarchical clustering. Hence the chosen value of  $K_g$ , and thus  $K$ , is preserved, which is the desired behavior.

In conclusion, let  $n_g$  denote the number of chain groups, then the initial mixture of output components for hierarchical clustering consists of  $K = K_g \cdot n_g$  components. Note that  $K_g$  is required input from the user, but  $n_g$  is determined automatically as a function of the critical  $R$  value  $R_c$ , a parameter that requires only moderate tuning. The initialization is described in pseudo code in Algorithm 1.

### 2.2.3. PMC run

The idea of adaptive importance sampling is to iteratively update the proposal function  $q$  to match the target density as closely as possible. In each step, regular importance sampling is carried out, and adaptation is performed with an expectation-maximization algorithm on mixture densities composed of Gaussian or Student’s  $t$  distributions [6, 7]. The proposal in step  $t$  is a mixture density

$$q^t(\boldsymbol{\theta}) = \sum_{j=1}^K \alpha_j^t q_j^t(\boldsymbol{\theta} | \boldsymbol{\xi}_j^t), \quad q_j^t \in \{\mathcal{N}, \mathcal{T}_\nu\}. \quad (6)$$

$q_j^t$  is a single multivariate component whose parameters are collectively denoted by  $\boldsymbol{\xi}_j^t = (\boldsymbol{\mu}_j^t, \boldsymbol{\Sigma}_j^t)$ . For the Student’s  $t$  case,  $q_j^t = \mathcal{T}_\nu$ ,  $\nu$  is the degree of freedom. The set of (normalized) component weights is denoted by  $\{\alpha_j^t : j = 1, \dots, K\}$ . Note that the component type,  $\mathcal{N}$  or  $\mathcal{T}_\nu$  (including  $\nu$ ), is fixed throughout a PMC run. The  $\mathcal{T}_\nu$  components may be preferable if the target has degeneracies or fat tails.

The goal in each update step is to reduce the Kullback-Leibler divergence  $\text{KL}(P \| q)$  towards the minimum value of 0 at  $q = P$ . The general problem of optimizing the KL functional is intractable. It is therefore necessary to reduce the complexity to an ordinary parameter optimization problem by fixing  $q$  to the form (6) and optimizing over  $\{(\alpha_j, \boldsymbol{\xi}_j) : j = 1, \dots, K\}$ . We want to remark that in the basic formulation of [7], the parameter  $\nu$  is held fixed, but it could be updated along with the other parameters through 1D numerical optimization [14].

For the Gaussian and Student’s  $t$  case, the updated values  $\alpha_j^{t+1}$  and  $\boldsymbol{\xi}_j^{t+1}$  are known, relatively simple-to-evaluate expressions [7] of  $q^t$  and the importance samples  $\{(\boldsymbol{\theta}_i^t, w_i^t) : i = 1 \dots N\}$  with importance weights  $w_i^t = P(\boldsymbol{\theta}_i^t)/q^t(\boldsymbol{\theta}_i^t)$ . It

---

**Algorithm 1** Initialization of the output components for hierarchical clustering. We use  $a = 0.2$  and  $R_c = 1.1 \dots 1.5$ .

---

**Require:**  $k$  chains with  $N_{\text{MCMC}}$  samples

**Require:** number of components per group  $K_g$

Start with empty initial mixture density  $q_{\text{HC}}^0(\cdot)$

Discard the first  $a \cdot N_{\text{MCMC}}$  samples for burn-in

Group the chains for given  $R_c$

$n_g \leftarrow$  number of groups

**for all** groups **do**

$k_g \leftarrow$  number of chains in group

**if**  $K_g < k_g$  **then**

        Merge chains into one long chain

$k_g \leftarrow 1$

**end if**

$n \leftarrow$  minimal lexicographic partition of  $K_g$  into  $k_g$  parts

**for all** chains in group **do**

        Partition into  $n_i$  patches

**for all** patches in chain **do**

            Compute sample mean  $\mu$  and covariance  $\Sigma$

            Add one component  $\mathcal{N}(\cdot | \mu, \Sigma)$  to  $q_{\text{HC}}^0(\cdot)$

**end for**

**end for**

**end for**

$q_{\text{HC}}^0(\cdot) \leftarrow$  Assign equal weight  $\frac{1}{n_g K_g}$  to every component

---

is important to stress again that PMC depends crucially on the initial proposal  $q^0$ , because the updates tend toward only a *local* minimum of KL.

Two useful quantities to determine when PMC updates become unnecessary because  $q$  is “close enough” to  $P$  are the *perplexity*  $\mathcal{P}$  [12] and the *effective sample size* ESS [24]. We normalize such that  $\mathcal{P}, \text{ESS} \in [0, 1]$ , where the optimal value of 1 is obtained for  $P \propto q$ , and thus  $w_i = w_j \forall i, j$ . While  $\mathcal{P}$  is sensitive rather to the mean of the distribution of the importance weights, ESS is a function of their variance. In case of successful PMC updates,  $\mathcal{P}$  rises monotonically until reaching a plateau (see for example [12, Fig. 6]). As discussed in detail in [5, Chapter 4.3], ESS is sensitive to outliers; i.e., samples with a weight much larger than the average. Outliers cause ESS to bounce up and down from one update to another (cf. [5, Fig. 4.7]), and render the ESS less robust as a convergence criterion. As these outliers seem to be inevitable in higher dimensions ( $d \gtrsim 25$ ), we focus on  $\mathcal{P}$  as the sole convergence criterion. We stop PMC updates when the relative increase of  $\mathcal{P}$  is less than 5% in consecutive steps. A value of  $\mathcal{P} > 0.9$  is attained only in simple low-dimensional problems. The PMC algorithm in abstract form, including our stoppage rule, is summarized in Algorithm 2.

The result of the first two stages of the algorithm, the MCMC prerun and the hierarchical clustering, is a Gaussian mixture density  $q_{\text{HC}}$ . Naïvely, we would set the initial proposal  $q^0 = q_{\text{HC}}$ , and start mapping the target density with PMC. However, a number of considerations have to be taken into account. We use Gaussians in the first two stages because the hierarchical clustering is then particularly fast and simple to implement. But we do not expect the chain patches turned into Gaussians to approximate the target density with the highest precision. In particular, many realistic problems have thicker tails, and are more accurately described by a Student’s  $t$  mixture. In fact, a much more involved hierarchical clustering for Student’s  $t$  exists [25], but we don’t expect it to reduce the number of PMC updates. The sole purpose of  $q_{\text{HC}}$  is to cover the support of the target with some accuracy, and the actual adaptation is left to the PMC update algorithm. In the end, we only use the samples drawn from the adapted PMC proposal for inference. We therefore consider it appropriate to perform two modifications to go from  $q_{\text{HC}}$  to  $q^0$ .

First, all component weights are set equal to balance the effect of an unequal number of chains in each group. The weights are adjusted properly in the first PMC update, so components are discarded if their target probability mass is low, and not because few chains visited them.

---

**Algorithm 2** The generic PMC algorithm. We use  $t_{\min} = 1, t_{\max} = 20, \varepsilon = 0.05$ .

---

**Require:** number of samples  $N$ , initial proposal  $q^0$

```

converged  $\leftarrow$  false,  $t \leftarrow 0$                                  $\triangleright$  Initialization
while  $(t < t_{\max}) \wedge (\neg \text{converged})$  do                         $\triangleright$  Update loop
    draw  $N$  samples  $\theta_i^t$  from  $q^t$  and compute importance weights  $w_i^t$ 
    if  $t \geq t_{\min} \wedge \left| \frac{p^t - p^{t-1}}{p^t} \right| < \varepsilon$  then
        converged  $\leftarrow$  true
    end if
     $q^{t+1} \leftarrow$  update proposal based on  $q^t$  and  $\{(\theta_i^t, w_i^t) : i = 1 \dots N\}$ 
     $t \leftarrow t + 1$ 
end while
if converged then                                                 $\triangleright$  Final step
    draw  $N_{\text{final}}$  samples from  $q_{\text{final}}$  and compute their importance weights
end if

```

---

Second, if a Student's t mixture is believed to yield a better representation of the target, we create a "clone" of  $q_{\text{HC}}$  where each Gaussian component is replaced by a Student's t component with identical location and scale parameter. The degree of freedom,  $\nu$ , is the same for all components, and currently has to be chosen a-priori by the user in the PMC approach. Its optimal value in the update is not known in closed form [7]. However, as noted in [14],  $\nu$  can be obtained from one-dimensional root finding. This is one source of future improvement, as guessing the proper value of  $\nu$  is not easy. In low dimensions, the difference is usually small, but for large  $d$ , the impact of outliers due to underestimating the tails may be significant, especially in plots of 2D marginal distributions (for an example plot, see [5, Fig. 4.10]).

Assuming that  $q^0$ , the initial proposal, is fixed, there is still an open question before we can start PMC: how many samples  $N$  to draw from the proposal? In the derivation of the PMC update step,  $N \rightarrow \infty$  is assumed, and this guarantees a reduced Kullback-Leibler divergence [7]. Large  $N$  ensures many samples from each component, but increases the computational burden. If  $N$  is too small, the updates may render  $q^{t+1}$  worse than  $q^t$ , and the PMC algorithm fails. A proper choice of  $N$  depends mostly on the dimensionality of the target density  $d$ ; for guidance, cf. the discussion in Section 3. After all,  $N$  is a required input to PMC, and is not deduced from the target. A reliable, quantitative rule to determine  $N$  would be very desirable, but is not available to us. We then attempt to ensure that every component is explored initially, so the quantity of interest is  $N_c$ , the *number of samples per component* in the first step, whence  $N = K \cdot N_c$ . Once the component weights are adjusted in the first update step, components that receive a very low relative weight are discarded, or "die"; i.e., the number of samples drawn from them is so small that there is not enough information gained to perform another update. In the reference implementation of PMC [11] that we use for the updating, the minimum number of samples per component is set at 20. We stop the update process when the convergence criteria of Algorithm 2 are met, and collect the samples used for inference in the final step. Note that we do not have to keep  $N$  constant in every step; in fact, we have experimented with reducing  $N$  as  $N = K_{\text{live}} \cdot N_c$ , where  $K_{\text{live}}$  is the number of live components. But we often saw PMC fail in those cases, as after a short number of steps,  $K_{\text{live}} \rightarrow 1$  resulting in  $\mathcal{P} \rightarrow 0$ . Therefore, we recommend using identical values  $N$  in every PMC update step for improved stability. For the accuracy of inference, a larger number of samples,  $N_{\text{final}}$ , is advisable in the final step. At any step,  $Z$  and its uncertainty are estimated from importance weights  $\{w_i : i = 1 \dots N\}$  through the sample mean and variance as

$$\widehat{Z} = \frac{1}{N} \sum_{i=1}^N w_i, \quad \widehat{V}[\widehat{Z}] = \frac{1}{N(N-1)} \sum_{i=1}^N (w_i - \widehat{Z})^2. \quad (7)$$

### 3. Short guide to parameter settings

At this point, we summarize the previous sections and provide guidance on setting the various tunable parameters. Crucial settings of the particular runs are listed in Table 1.

For the MCMC step, we use  $k = 10 \dots 50$  chains depending on the expected number of target modes. For a simple unimodal distribution, a handful of chains should suffice. The chains are run for  $N_{\text{MCMC}} = 10\,000$  ( $d \gtrsim 2$ ) –  $100\,000$  ( $d = 42$ ) iterations with a Gaussian proposal, though Student’s  $t$  could be used as well. For simple problems with  $d$  small and a decent initial chain proposal, the minimum value of  $N_{\text{MCMC}}$  is on the order of 1000. Discarding the initial 20 % for burn-in, we split up the chains into patches of length  $L = 50 - 300$ , the exact value of  $L$  is not critical.

With regard to hierarchical clustering, we group chains according to the  $R$  values, using a threshold value around  $R_c = 1.2$ . For larger dimensions or respectively smaller  $N_{\text{MCMC}}$ , larger values up to 1.5 or even 2 can be used. The number of components per group,  $K_g$ , ought to be  $\gtrsim d$ ; the bigger  $K_g$ , the more accuracy is obtained at the expense of more evaluations of the target. The initial components arise from long patches of chains within a group. Hierarchical clustering is stopped if the distance measure in two consecutive steps is reduced by less than  $\varepsilon_{\min} = 10^{-4}$ .

In the PMC step, we initially set all component weights equal. In most applications, a Gaussian mixture has tails that are thinner than the target’s tails, so one can decide for a Student’s  $t$  mixture with degree of freedom  $\nu = 2 \dots 15$ . Good results were obtained with  $N_c$  ranging from 200 ( $d = 2$ ) over 600 ( $d = 20$ ) to 2500 ( $d = 42$ ). Convergence is declared when the normalized perplexity  $\mathcal{P}$  is stable to within 5 % between two consecutive steps. Jumps in the ESS hint at outliers caused by too few mixture components or by a proposal whose tails are too thin. If the PMC updates “kill” more and more components and reduce the perplexity, more initial components and a larger sample size may help. Another improvement may be to slightly increase  $N_{\text{MCMC}}$  or  $k$ . If outliers have a dominant effect on the resulting marginal distributions, the combined effect of smoothing with kernel density estimation and outlier removal provides a partial remedy. After convergence, a final sample size  $N_{\text{final}}$  of as small as 5000 is sufficient for an integral estimate at the percent level when  $\mathcal{P} \lesssim 1$  and  $d$  small. For the toughest problems where  $\mathcal{P}$  remains low, a size ranging in the millions is necessary for targets in  $d \gtrsim 30$ .

## 4. Examples

As described in the introduction, few publicly available codes exist for the solution of difficult analysis problems. One is Multinest [17], and we use this to benchmark our new approach in the following examples.

### 4.1. Gaussian shells

For easy comparison with the Multinest package, we use the same Gaussian-shell example discussed in [17, Sec. 6.2] of two well separated hyperspheres with a Gaussian density profile. We define the likelihood as

$$L(\boldsymbol{\theta}) = \frac{1}{2} \text{circ}(\boldsymbol{\theta}|\mathbf{c}_1, r, w) + \frac{1}{2} \text{circ}(\boldsymbol{\theta}|\mathbf{c}_2, r, w), \quad \text{circ}(\boldsymbol{\theta}|\mathbf{c}, r, w) = \frac{1}{\sqrt{2\pi}w^2} \exp\left[-\frac{(|\boldsymbol{\theta} - \mathbf{c}| - r)^2}{2w^2}\right]. \quad (8)$$

The width  $w = 0.1$  is chosen small compared to the radius  $r = 2$ , emulating a problem with a continuous spherical symmetry (degeneracy). The two shells are well separated by a distance of 7 units and do not overlap. The shell centers are placed at

$$\mathbf{c}_{1,2} = (\pm 3.5, 0, \dots, 0), \quad (9)$$

and uniform priors over a hypercube  $\boldsymbol{\theta} \in [-6, 6]^d$  are assumed. An accurate analytical approximation of the evidence is

$$Z = \frac{\sqrt{2\pi}^{(d-1)/2}}{\Gamma(d/2)(2\rho_{\max})^d w} \int_0^{\rho_{\max}} d\rho \rho^{d-1} \exp\left(-\frac{(\rho - r)^2}{2w^2}\right), \quad (10)$$

where the integral is performed over a hypersphere of radius  $\rho_{\max} = 6$  covering a single shell. For the case at hand, the contribution from the likelihood in the region contained in the hypercube but not in the hypersphere is negligible. Note that there is an extra factor of 1/2 in our definition (8) compared to [17].

To assess the algorithm’s performance, we repeat the run 100 times with different pseudo-random numbers in  $d = 2, 10, 20$  dimensions. The parameter settings are listed in Table 1. For comparison, we also run Multinest 100 times with parameter settings as advocated in [17] with 1000 live points and desired sampling acceptance rate of 0.3. To allow an easier comparison, we fix the number of samples in the final PMC step,  $N_{\text{final}}$ , at the average number of samples that Multinest yields. Note that this does not equal the total number of target evaluations,  $N_{\text{total}}$ , with either



algorithm. The Multinest algorithm accepts samples only with a certain rate  $\varepsilon$  such that  $N_{\text{total}} = N_{\text{final}}/\varepsilon$ . For our algorithm, the samples at the MCMC stage and during the  $t_{\text{final}}$  PMC updates have to be added such that

$$N_{\text{total}} = kN_{\text{MCMC}} + t_{\text{final}}KN_c + N_{\text{final}}. \quad (11)$$

### Discussion

We now comment on the performance of the algorithm in  $d = 2$ , with relevant settings and results summarized in Table 1 and Table 2. In the MCMC step, individual chains may or may not explore the entire region of one shell. Examples are shown in the top left panel of Fig. 2. It is important that both shells are discovered, and that the combination of chains covers both regions (Fig. 2, top right), although the relative masses of the shells are incorrect because more chains visit the left shell (5 versus 3). The large number of 640 chain patches is used in hierarchical clustering to convert the initial guess with 30 components (Fig. 2, center left) into  $q_{\text{HC}}$  (Fig. 2, center right). The fact that approximately one third of the components present in the initial guess are discarded during the clustering demonstrates that the clustering may detect an unnecessary surplus of components, a welcome feature. Nonetheless, there is a small number of “outlier” components in  $q_{\text{HC}}$ , distinguished by the larger size and location in the interior of a shell. These components are assigned a vanishing weight during the 5 PMC updates needed to obtain the final proposal  $q_{\text{final}}$  (Fig. 2, bottom left). The final result (Fig. 2, bottom right) accurately captures the two shells and assigns equal probability mass to either shell, in contrast to the results obtained from the combination of chains

In all of the 100 repetitions, the initialization is successful. PMC converges quickly and determines the evidence to an accuracy of roughly 1 % from only 5200 samples in the final step. Perplexity and ESS take on large values, implying a good approximation of the target by the mixture density. Note that the number of “active” components in the final step at about  $K_{\text{final}} = 17$  is significantly lower than the  $K_g N_g = 30$  components available at the beginning of hierarchical clustering. During the clustering, on average 10 components are discarded, and only about 3 become inactive during the PMC updates. Extending to  $d = 10$  and 20, the evidence is again determined accurately at the percent level. We notice that the fraction of runs,  $f$ , in which the correct evidence is contained in  $[\widehat{Z} - \widehat{\Delta Z}, \widehat{Z} + \widehat{\Delta Z}]$  diminishes slightly as  $d$  increases, but remains at a very reasonable value of  $f = 61$  % in  $d = 20$ .

For  $d > 2$ , somewhat surprisingly fewer PMC updates ( $t_{\text{final}} = 2 - 3$ ) are needed than in two dimensions ( $t_{\text{final}} \approx 5$ ). This is likely due to having relatively more components, and thus flexibility, available in the proposal in  $d = 2$ . For  $d > 2$ ,  $\mathcal{P}$  and ESS settle to lower values around 30–40 %. This reduction in the maximum attainable  $\mathcal{P}$  is common in higher dimensions—a mild form of the “curse of dimensionality”. We note that both shells are discovered, and all proposal components remain active throughout the clustering and PMC updates. This demonstrates a successful adaptation due to the good initialization from chains and hierarchical clustering. By graphical inspection, we verified that marginal distributions agree well with expectations.

In the following, we want to compare directly how PMC and Multinest perform. In general, both algorithms perform well; the two shells are properly explored, and the correct evidence is found in the estimated interval in roughly 2/3 of the runs (see Fig. 3 and Table 2). The uncertainty estimate  $\widehat{\Delta Z}$  is 5–20 times smaller with PMC for the same number of samples considered, but at the expense of a larger  $N_{\text{total}}$ . We list the total running time and the average number of calls in Table 3. To ensure a fair comparison, we did not use any parallel evaluations, even though this is one of PMC’s main strengths. Furthermore, we replaced Multinest’s slow default output to text files by output to the binary HDF5 format [26]—the same format we use to store chains and PMC output. In this example, the target density is quick to evaluate, hence the majority of time is spent on updating the proposal in the PMC case, or on clustering and finding a new point in Multinest.

As shown in Table 3, Multinest requires less times to run in all considered cases and exhibits lower  $N_{\text{total}}$ . While the exact numbers depend on the algorithms’ parameter settings, we observe that Multinest is ahead on simpler problems, but our algorithm becomes competitive in higher dimensions as Multinest’s acceptance rate drops and  $N_{\text{total}}$  increases only moderately with  $d$  for our algorithm. We repeat that the timings were done in serial execution, but PMC can easily use a large number of cores simultaneously and hence is preferable for very costly target functions.

When it comes to integration precision, PMC is to be preferred, as it is able to determine the evidence at an accuracy of better than one percent even in  $d = 20$ . In case the precision is considered too low *after* the final step, it is straightforward to load the stored final proposal from disk to sample more points until reaching the desired precision with the usual  $1/\sqrt{N_{\text{final}}}$  scaling. In contrast, Multinest’s precision is significantly lower, even by a factor of roughly 20 in  $d = 20$ , and cannot simply be improved by continuing the run because of its serial nature of samples ordered

by likelihood value. For both algorithms, the relative estimated uncertainty on  $Z$  averaged over the 100 repetitions,  $E\left[\frac{\widehat{\Delta Z}}{\widehat{Z}}\right]$ , agrees well with the relative spread of the distribution of evidence estimates,  $\frac{\sigma[\widehat{Z}]}{E[\widehat{Z}]}$ . The agreement is good to the last significant digit given in Table 2 for PMC, confirming the usefulness of the uncertainty estimate and matching the near-Gaussian shape of the distributions of  $\widehat{Z}$  shown in Fig. 3.

	$d$	$k$	$N_{\text{MCMC}}$	$L$	$\nu$	$K_g$	$N_c$	$N_{\text{final}}$	$K_{\text{final}}$	$t_{\text{final}}$	$\mathcal{P}/\%$	ESS /%
Shells	2	8	10 000	100	-	15	200	5200	17.3	5.02	75	51
	10	8	20 000	100	-	15	400	18 000	30	2	35	30
	20	8	20 000	200	-	25	600	40 000	50	2.9	43	36
Tails	2	20	10 000	100	12	5	200	6700	19.95	1.4	95	93
	10	20	20 000	100	12	15	400	30 000	58.8	2.2	81	71
	20	20	20 000	200	12	25	600	54 000	99.91	3.5	56	34

Table 1: Settings and results of MCMC + PMC runs for the examples in  $d$  dimensions. Each of the  $k$  chains is run for  $N_{\text{MCMC}}$  iterations split into patches of length  $L$ .  $\nu$  is the degree of freedom of each component in each  $\mathcal{T}$  mixture proposal density, a missing value represents a  $\mathcal{N}$  mixture.  $K_g$  is the number of mixture components per group of chains, and  $N_c$  is the number of samples per component drawn during the first PMC update step, while  $N_{\text{final}}$  is the fixed number of samples from all components in the final step.  $K_{\text{final}}$  is the number of active components after  $t_{\text{final}}$  PMC updates in the final step, in which  $\mathcal{P}$  and ESS characterize the quality of the adaptation of the proposal to the target.  $K_{\text{final}}$ ,  $t_{\text{final}}$ ,  $\mathcal{P}$ , and ESS are averaged over 100 runs with the following common settings. During MCMC, the Gaussian local random walk proposal function is updated after  $N_{\text{update}} = 200$  iterations in  $d = 2$  and  $N_{\text{update}} = 500$  iterations in  $d > 2$ . Chains are grouped according to a critical  $R$  value of  $R_c = 1.2$ , and the first 20 % iterations are discarded for burn-in.

		PMC					Multinest				
	$d$	$Z$	$E[\widehat{Z}]$	$\frac{\sigma(\widehat{Z})}{E[\widehat{Z}]}$	$E\left[\frac{\widehat{\Delta Z}}{\widehat{Z}}\right]$	$f/\%$	$E[\widehat{Z}]$	$\frac{\sigma(\widehat{Z})}{E[\widehat{Z}]}$	$E\left[\frac{\widehat{\Delta Z}}{\widehat{Z}}\right]$	$f/\%$	
Shells	2	$8.726 \cdot 10^{-2}$	$8.73 \cdot 10^{-2}$	0.008	0.009	69	$8.9 \cdot 10^{-2}$	0.06	$^{+0.052}_{-0.050}$	57	
	10	$2.304 \cdot 10^{-7}$	$2.30 \cdot 10^{-7}$	0.011	0.012	71	$2.4 \cdot 10^{-7}$	0.13	$^{+0.13}_{-0.12}$	68	
	20	$1.064 \cdot 10^{-16}$	$1.06 \cdot 10^{-16}$	0.007	0.007	61	$1.1 \cdot 10^{-16}$	0.15	$^{+0.21}_{-0.17}$	80	
Tails	2	$2.778 \cdot 10^{-4}$	$2.78 \cdot 10^{-4}$	0.003	0.003	72	$2.8 \cdot 10^{-4}$	0.07	$^{+0.064}_{-0.060}$	63	
	10	$1.654 \cdot 10^{-18}$	$1.66 \cdot 10^{-18}$	0.004	0.004	63	$6.3 \cdot 10^{-18}$	0.13	$^{+0.17}_{-0.14}$	0	
	20	$2.735 \cdot 10^{-36}$	$2.73 \cdot 10^{-36}$	0.006	0.006	61	$3.3 \cdot 10^{-34}$	0.37	$^{+0.24}_{-0.19}$	0	

Table 2: Performance metrics of the 100 runs for the examples in  $d$  dimensions.  $Z$  is the true evidence value.  $E[\cdot]$  and  $\sigma[\cdot]$  denote the sample mean and respectively the square root of the sample variance across all runs, whereas  $\widehat{Z}$  and  $\widehat{\Delta Z}$  denote the evidence and respective uncertainty estimate from a single run according to (7).  $f$  is the fraction of runs in which  $Z$  is contained in  $[\widehat{Z} - \widehat{\Delta Z}, \widehat{Z} + \widehat{\Delta Z}]$ . For the tail data in  $d = 2, 10$ , we use only the runs covering all four modes. Multinest’s uncertainty estimate is transformed from the log scale to the linear scale and thus becomes asymmetric.

#### 4.2. Heavy tails

The second example has four well separated maxima, each with the same shape and probability mass. Such a distribution arises naturally in many high-energy physics analyses in which the underlying physics model has discrete symmetries in the parameters of interest. If the symmetry is not exact, it is of great interest to accurately determine the relative masses of the maxima. We define individual maxima as products of Gaussian and LogGamma 1D distributions. The LogGamma distribution [27] is an asymmetric heavy-tailed distribution with a location, scale, and shape parameter. In this form, the example is a simplified yet hard-to-solve version of the posterior appearing in our motivating analysis [4].

For simplicity, we begin in  $d = 2$  dimensions where  $L(\theta) = L(\theta_1, \theta_2) = L(\theta_1) \cdot L(\theta_2)$  (see top left panel of Fig. 4).  $L(\theta_1)$  is a mixture of two LogGamma components with maxima at  $\theta_1 = \pm 10$  with unit scale and unit shape parameter.

	$d$	MCMC+PMC		Multinest		$\varepsilon/\%$
		$t/s$	$N_{\text{total}}$	$t/s$	$N_{\text{total}}$	
Shells	2	1.6	105 000	0.95	8300	63
	10	5.4	202 000	2.9	50 000	36
	20	45	274 000	12.3	208 000	19
Tails	2	2.9	212 300	0.4	18 000	38
	10	23	482 800	13.7	269 000	10
	20	166	628 000	108	5 091 000	1

Table 3: Typical values of run time  $t$ , total number of target density evaluations  $N_{\text{total}}$ , and acceptance rate  $\varepsilon$  (Multinest only) solving the Gaussian shell (upper half) and heavy-tailed mode (lower half) examples on a single core of an Intel i7 920 clocked at 2.67 GHz.

Similarly,  $L(\theta_2)$  is a mixture of standard normal distributions centered around  $\theta_2 = \pm 10$ :

$$\begin{aligned} L(\theta_1) &= 0.5 \text{LogGamma}(\theta_1|10, 1, 1) + 0.5 \text{LogGamma}(\theta_1| - 10, 1, 1) \\ L(\theta_2) &= 0.5 \mathcal{N}(\theta_2|10, 1) + 0.5 \mathcal{N}(\theta_2| - 10, 1). \end{aligned} \quad (12)$$

To explore the effect of higher dimensions, we augment the 2D target density with an equal number of LogGamma and Gaussian distributions. Throughout, we assume uniform priors on  $\theta \in [-30, 30]^d$ . Specifically, the target likelihood is

$$L(\theta) = \prod_{i=1}^d L(\theta_i), \quad (13)$$

where the distribution in the first two dimensions is given by (12), and

$$L(\theta_i) = \begin{cases} \text{LogGamma}(\theta_i|10, 1, 1), & 3 \leq i \leq \frac{d+2}{2} \\ \mathcal{N}(\theta_i|10, 1), & \frac{d+2}{2} < i \leq d. \end{cases} \quad (14)$$

Since  $L(\theta)$  is normalized to unity, the evidence in  $d$  dimensions is given by the prior normalization as  $Z(d) = 60^{-d}$ . There are four modes due to the first two dimensions; the extra dimensions do not add any further modes. Hence, we perform  $R$ -value grouping as well as clustering in Multinest only in the first two dimensions. The settings for MCMC and the initialization are summarized in the lower left half of Table 1. In contrast to the Gaussian-shell example, we now use more chains ( $k = 20$ ) to increase the chances to discover all modes, vary  $K_g$  more strongly with  $d$ , and choose a  $\mathcal{T}$  mixture with  $\nu = 12$  degrees of freedom in response to the heavy tails of LogGamma.

### Discussion

As with the Gaussian shells, we repeat the analysis 100 times with both MCMC+PMC and Multinest. In the lower half of Table 1, we list the convergence properties of our algorithm. In all probed dimensions, essentially all components remain active, convergence is reached after less than five steps even in  $d = 20$ , and perplexity reaches high levels. In summary, the initialization procedure works well.

The results of the evidence calculations are shown in Table 2 and Fig. 3. For PMC, both the estimated and the actual accuracy are around an excellent value of 0.5 %. Note a common pitfall of the Markov chain approach apparent in the bimodal distribution for  $d = 2$  and  $d = 10$  in Fig. 3. In one ( $d = 2$ ), respectively eight ( $d = 10$ ), of the runs, only three of the four modes are discovered, hence the evidence is off by 25 % and there is no way for the algorithm to know something is missing. In an actual data analysis, one could gain extra knowledge about the target, say from repeated mode finding, to learn about the existence and location of various modes. Seeding the chains in different modes would make the MCMC step both more reliable and more efficient. However, we use this example both to show how well the algorithm performs with no such extra knowledge given—just a fairly large number of  $k = 20$  chains—and to alert unsuspecting users. Future developments should remedy this issue in a more robust fashion; some ideas are discussed in Section 5. Focusing on the runs in which all four modes are found, the fraction  $f$  of runs for which the true value is in  $[\widehat{Z} - \widehat{\Delta Z}, \widehat{Z} + \widehat{\Delta Z}]$  gradually decreases to 61 % in  $d = 20$ .

Marginal distributions estimated from the PMC accurately approximate the target density as shown in Fig. 4, where we show 1D and 2D marginals in  $d = 20$  for one example run. For  $d \gtrsim 20$ , one starts to observe outliers. A mild outlier is visible in Fig. 4 near  $(\theta_1, \theta_2) = (2, 11)$ . Outliers are less of a problem in 1D marginals; an example is shown in logarithmic and linear scale in Fig. 4. However, we stress that it is a common problem that a single outlier can dominate a 2D marginal in  $d \gtrsim 30$  if the proposal density does not perfectly match the target. This is another manifestation of the curse of dimensionality that plagues importance sampling methods in general. Some remedies include smoothing of the marginal and filtering outliers [5].

For comparison, we also ran Multinest with 1000 live points and the desired acceptance rate set to 0.3 in all dimensions; cf. [17] for a description of Multinest’s parameters. Regarding mode finding, Multinest is very robust as it discovers all four modes in every run. In  $d = 2$ , Multinest performs reasonably well, with an acceptance rate of 38 % and estimated accuracy of 7 % with about a factor of 10 fewer calls to the target. For  $d = 10$ , Multinest’s acceptance rate reduces to 10 %, and the evidence is overestimated<sup>1</sup> by at least a factor of 2.5 in all runs, the coverage dropping to zero. Similarly in  $d = 20$ , the evidence estimate is too large by a factor of at least 40, with  $f = 0$  despite an estimated uncertainty of roughly 20 %. The probability mass of an individual mode ranges from 10 – 58 % compared to the correct value 25 %. However, the 1D marginals agree well with expectations, and furthermore there is no problem with outliers in Multinest, presumably because the weight associated with each sample is based on a stochastic estimate of the prior mass and is thus intrinsically smoothed.

The CPU usage is listed in Table 3. While Multinest roughly needs a factor of 13 *fewer* target evaluations in  $d = 2$ , it requires eight times *more* in  $d = 20$  due its low acceptance rate of 1 %. Nonetheless, PMC takes about 60 % longer for this simple target; this extra times is spent almost exclusively in the proposal updates due the large number of 100 components. Ideas for faster updates are discussed below.

## 5. Outlook

The proposed initialization of PMC with Markov chains and hierarchical clustering performs well in the above examples. However, there are still numerous improvements to make en route to the ideal black-box sampler.

### Initialization

The two most important aspects of the algorithm design to us are *correctness* and *speed*. Regarding the former, it is crucial to ensure that the algorithm leads to samples from the target; i.e., we need to improve the automatic detection of *all* regions of parameter space that contribute. Within our framework, this could be achieved by choosing the chains’ starting points more cleverly, perhaps based on a preliminary sampling step. This should lead to a reduction of the necessary number of chains,  $k$ , for the majority of problems that have at most a handful of separated maxima.

Concerning speed, we consider reducing the overall execution time, and also reducing the number of parameters steering the algorithm. The latter helps in two ways. First, it requires time to assign a good value to each parameter, either at run-time or by the user performing repeated trials. Second, a user may inadvertently make a poor choice with adverse effects on the performance. The most important parameter in this regard is  $K$ , the number of mixture components. A promising alternative to hierarchical clustering is the *variational Bayes* approach described in [28], in which the “best” number of components is computed along with the positions and covariances of the reduced mixture’s components. We could obviate the patch length  $L$  by giving up slicing the chains into patches if we instead did the clustering at the level of individual chain samples.

More fundamentally, one could eliminate the MCMC prerun entirely in favor of a large number of samples from the prior or a uniform distribution on the parameter space. Two advantages are that the samples can be computed with massive parallelization and that potentially fewer samples are required by avoiding the redundancy of multiple chains in the same region. This path is followed in [15] and shown to work for unimodal problems up to  $d = 20$ . But by giving up the MCMC prerun, we suspect there is a greater chance that suppressed modes and degeneracies are missed or poorly captured in high-dimensional problems. In addition, it proved useful for validation purposes to compare the marginal distributions from MCMC and PMC for qualitative agreement. If a region is visible in the MCMC but not

---

<sup>1</sup>During the final stages of preparing this article, we were made aware of an ongoing effort by the Multinest authors to significantly improve the integration accuracy in the upcoming version 3.0.

in the PMC output, either PMC failed, or that region contains negligible probability mass, which can be verified with the samples' target values in that region. We expect only minor improvements when replacing the Gaussian clustering with the considerably more involved Student's  $t$  clustering to obtain a Student's  $t$  mixture proposal density from the chain patches [25].

### PMC

Apart from the initialization, there are more general directions to further enhance PMC. Our examples suggest that, given a good initial proposal, importance sampling works well up to  $d \approx 30$ , and in fact we successfully sampled from a single mode of the heavy-tailed example in  $d = 42$ , but problems with outliers appear already for  $d \gtrsim 20$ . To a certain degree, outliers can be reduced by more mixture components, an adjustment of the Student's  $t$  degree of freedom  $\nu$ , and more samples per component. We presented guidance how to manually adjust these parameters, but an automatic adjustment is highly preferred.  $\nu$  can be determined in each PMC update by a 1D numerical solution of Eq. (16) in [14], making an informed user guess for  $\nu$  obsolete. Providing even larger flexibility, individual components may then even have different values of  $\nu$ . The soft limit of  $d \approx 40$  ([14] noted a maximum of  $d = 35$  in their applications) is due not only to outliers but also to the excessive time of updating the proposal, whereas the MCMC initialization still works well.

It is conceivable that marginal distributions are less affected by outliers if the proposal function of the final PMC step,  $q_{\text{final}}$ , is used as a global proposal in MCMC. Using  $q_{\text{final}}$  as a global proposal alone would not solve that issue because the Metropolis-Hastings acceptance probability [29] to accept a new point  $\theta_2$  given the current point  $\theta_1$  is just the ratio of importance weights  $w_2/w_1$ . An outlier would be accepted with probability close to one, and the chain would then be stuck for many iterations. Hence one needs a mixture of global and local jumps for efficient sampling. We envision that—after appropriate scaling of the covariance matrices—individual mixture components guide local jumps, and the full proposal is used for global jumps. Using a fixed proposal and assuming rapid mixing due to the global jumps, massive parallelization is straightforward and individual chains need to be run for fairly few iterations. Similar efforts, though still involving a fairly large number of ad-hoc choices, are reported in [30].

The curse of dimensionality surfaces because the PMC update scales as  $O(KN_c d^4)$ , where  $K$  and  $N_c$  have to be chosen larger for increasing  $d$ . Due to the Rao-Blackwellization [7], each component density is evaluated for every one of the  $KN_c$  samples, and each such evaluation requires a vector  $\times$  matrix  $\times$  vector product. Currently the update is executed in serial code, but could be massively parallelized, and is easily computed together with the target density. Another way to speed up would be to partition the components such that for samples from the  $j^{\text{th}}$  component only the subset of other components needs to be considered that has a substantial contribution. Ideally, the partitioning—similar to that employed in an implementation of the fast Gauss transform [31]—could be done at the component level before the weights are computed so each parallel process would do only the minimum required evaluations.

## 6. Conclusion

A new method is introduced to solve initialization difficulties in adaptive importance sampling. The method was initially developed in the context of a global fit to large data sets [4, 5]. It uses a combination of Markov chains, hierarchical clustering [18], and population Monte Carlo [6, 7]. Our method was compared to a publicly available implementation of nested sampling [17] for examples with multimodal posterior distributions in up to 20 dimensions, and found to perform well. The evidence was more accurate than that from nested sampling and the marginal distributions were found to reproduce the target distribution. The algorithm is amenable to massive parallelization.

The main development of this work consists in providing a reliable initialization of adaptive importance sampling, allowing it to converge in very few steps. While some tuning of parameters is still necessary, we consider this to be an important step toward a “black-box” sampling algorithm.

## Acknowledgments

The authors thank Danny van Dyk and Christoph Bobeth for numerous inspiring discussions during the early development of the algorithm. We are grateful to Martin Kilbinger for hints on using his PMCLib package, and to Farhan Feroz for guidance on Multinest. F. Beaujean wishes to thank the IMPRS for elementary particle physics for the generous support of this work.

## References

- [1] C. Robert, G. Casella, Monte Carlo statistical methods, Springer, 2004.
- [2] S. S. AbdusSalam, B. C. Allanach, F. Quevedo, F. Feroz, M. Hobson, Fitting the phenomenological MSSM, Phys. Rev. D 81 (2010) 095012.
- [3] P. Ade, et al., Planck 2013 results. XVI. Cosmological parameters, Submitted to Astronomy & Astrophysics. [arXiv:1303.5076](#).
- [4] F. Beaujean, C. Bobeth, D. van Dyk, C. Wacker, Bayesian Fit of Exclusive  $b \rightarrow s\bar{\ell}\ell$  Decays: The Standard Model Operator Basis, JHEP 1208 (2012) 030.
- [5] F. Beaujean, A Bayesian analysis of rare B decays with advanced Monte Carlo methods, Dissertation, Technische Universität München (2012).  
URL <http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20121114-1115832-1-8>
- [6] O. Cappé, A. Guillin, J.-M. Marin, C. P. Robert, Population Monte Carlo, J. Comput. Graph. Statist. 13 (4) (2004) 907–929.
- [7] O. Cappé, R. Douc, A. Guillin, J.-M. Marin, C. P. Robert, Adaptive importance sampling in general mixture classes, Stat. Comp. 18 (2008) 447–459.
- [8] J. von Neumann, N. Metropolis, S. Ulam, Monte Carlo Method, National Bureau of Standards/Applied Math. Series 12 (1951) 36–38.
- [9] A. Dempster, N. Laird, D. Rubin, Maximum likelihood from incomplete data via the EM algorithm, J. R. Stat. Soc. Ser. B Stat. Methodol. 39 (1) (1977) 1–38.
- [10] S. Kullback, R. Leibler, On information and sufficiency, Ann. Math. Stat. 22 (1) (1951) 79–86.
- [11] M. Kilbinger, K. Benabed, O. Cappé, J. Coupon, J.-F. Cardoso, G. Fort, H. J. McCracken, S. Prunet, C. P. Robert, D. Wraith, PMC lib v1.0 (2011).  
URL <http://www2.iap.fr/users/kilbinge/CosmoPMC/>
- [12] D. Wraith, M. Kilbinger, K. Benabed, O. Cappé, J.-F. Cardoso, et al., Estimation of cosmological parameters using adaptive importance sampling, Phys. Rev. D 80 (2009) 023507.
- [13] M. Kilbinger, D. Wraith, C. P. Robert, K. Benabed, O. Cappe, et al., Bayesian model comparison in cosmology with Population Monte Carlo, Mon. Not. R. Astron. Soc. 405 (4) (2010) 2381–2390.
- [14] L. Hoogerheide, A. Opschoor, H. K. van Dijk, A class of adaptive importance sampling weighted EM algorithms for efficient and robust posterior and predictive simulation, J. Econometrics 171 (2) (2012) 101 – 120.
- [15] J.-M. Cornuet, J.-M. Marin, A. Mira, C. P. Robert, Adaptive Multiple Importance Sampling, Scand. J. Stat. 39 (4) (2012) 798–812.
- [16] J. Skilling, Nested sampling for general Bayesian computation, Bayesian Analysis 1 (4) (2006) 833–860.
- [17] F. Feroz, M. Hobson, M. Bridges, MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics, Mon. Not. Roy. Astron. Soc. 398 (2009) 1601–1614.
- [18] J. Goldberger, S. Roweis, Hierarchical clustering of a mixture model, Adv. Neur. Info. Proc. Syst. 17 (2004) 505.
- [19] H. Haario, E. Saksman, J. Tamminen, An Adaptive Metropolis Algorithm, Bernoulli 7 (2) (2001) pp. 223–242.
- [20] G. O. Roberts, A. Gelman, W. R. Gilks, Weak Convergence and Optimal Scaling of Random Walk Metropolis Algorithms, Ann. Appl. Probab. 7 (1) (1997) pp. 110–120.
- [21] G. Schwarz, Estimating the dimension of a model, Ann. Stat. 6 (2) (1978) 461–464.
- [22] H. Akaike, A new look at the statistical model identification, IEEE Transactions on Automatic Control 19 (6) (1974) 716–723.
- [23] A. Gelman, D. Rubin, Inference from iterative simulation using multiple sequences, Stat. Sci. 7 (4) (1992) 457–472.
- [24] J. Liu, R. Chen, Blind deconvolution via sequential imputations, J. Amer. Statist. Assoc. 90 (430) (1995) 567–576.
- [25] A. El Attar, A. Pigeau, M. Gelgon, Fast aggregation of Student mixture models, in: European Signal Processing Conference (Eusipco 2009), 2009, pp. 312–216.
- [26] The HDF5 group, Hierarchical data format version 5 (2013).  
URL <http://www.hdfgroup.org/HDF5>
- [27] G. Crooks, The Amoroso distribution (2010). [arXiv:1005.3274](#).
- [28] P. Bruneau, M. Gelgon, F. Picarougne, Parsimonious reduction of Gaussian mixture models with a variational-Bayes approach, Patt. Recog. 43 (3) (2010) 850–858.
- [29] W. Hastings, Monte Carlo sampling methods using Markov chains and their applications, Biometrika 57 (1) (1970) 97–109.
- [30] P. Giordani, R. Kohn, Adaptive Independent Metropolis–Hastings by Fast Estimation of Mixtures of Normals, J. Comput. Graph. Statist. 19 (2) (2010) 243–259.
- [31] V. I. Morariu, B. V. Srinivasan, V. C. Raykar, R. Duraiswami, L. S. Davis, Automatic online tuning for fast Gaussian summation, Adv. Neural. Info. Proc. Syst 21 (2009) 1113–1120.

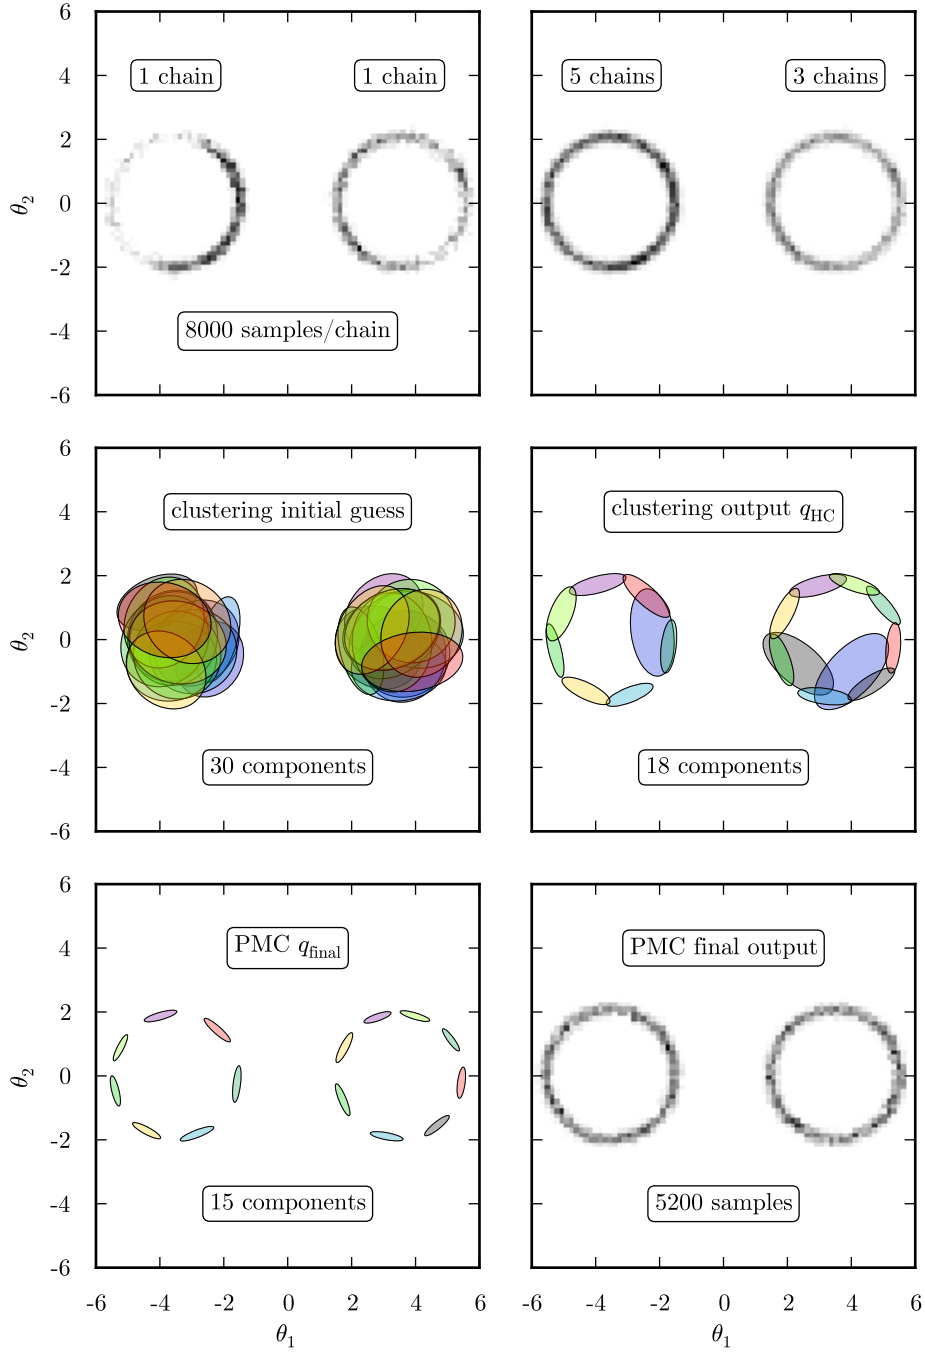


Figure 2: Gaussian shell example in  $d = 2$ . Histogram estimate of  $P(\theta_1, \theta_2)$  from single chains (top left) and from combining all  $k = 8$  chains (top right). 1- $\sigma$  contours of the individual components of the Gaussian mixture density in hierarchical clustering for the initial guess (center left) and output (center right) and in the final PMC step (bottom left). The same color code is used in the latter two figures to identify components. Histogram estimate of  $P(\theta_1, \theta_2)$  with samples of the final PMC step (bottom right).

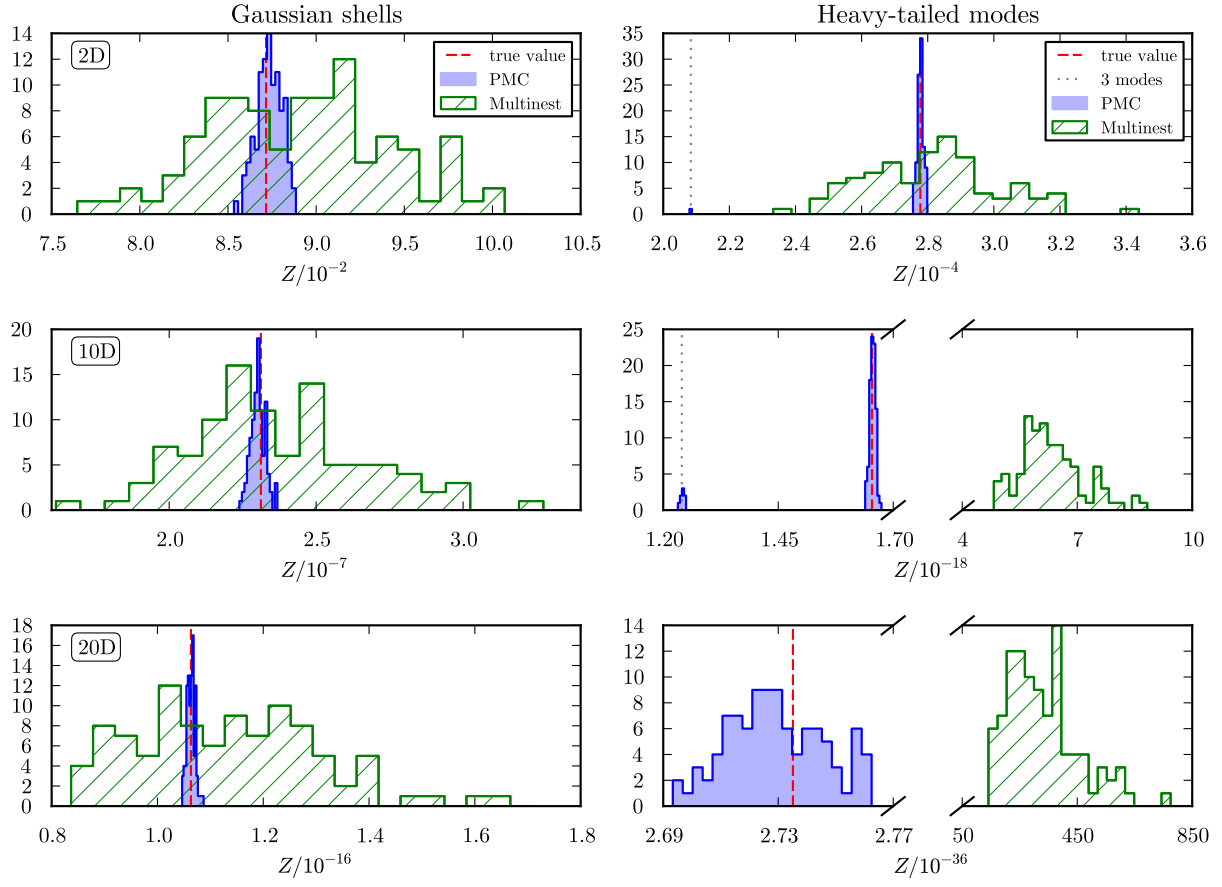


Figure 3: Evidence estimates of 100 runs of PMC with MCMC initialization and Multinest for the Gaussian shell (left column) and heavy tail (right column) examples in  $d = 2$  (top row),  $d = 10$  (center row), and  $d = 20$  (bottom row) dimensions. Note the very different scales used in each of the two lower plots on the right-hand side. The dashed vertical line indicates the true evidence, whereas the dotted line gives the evidence when one of the modes is missed.



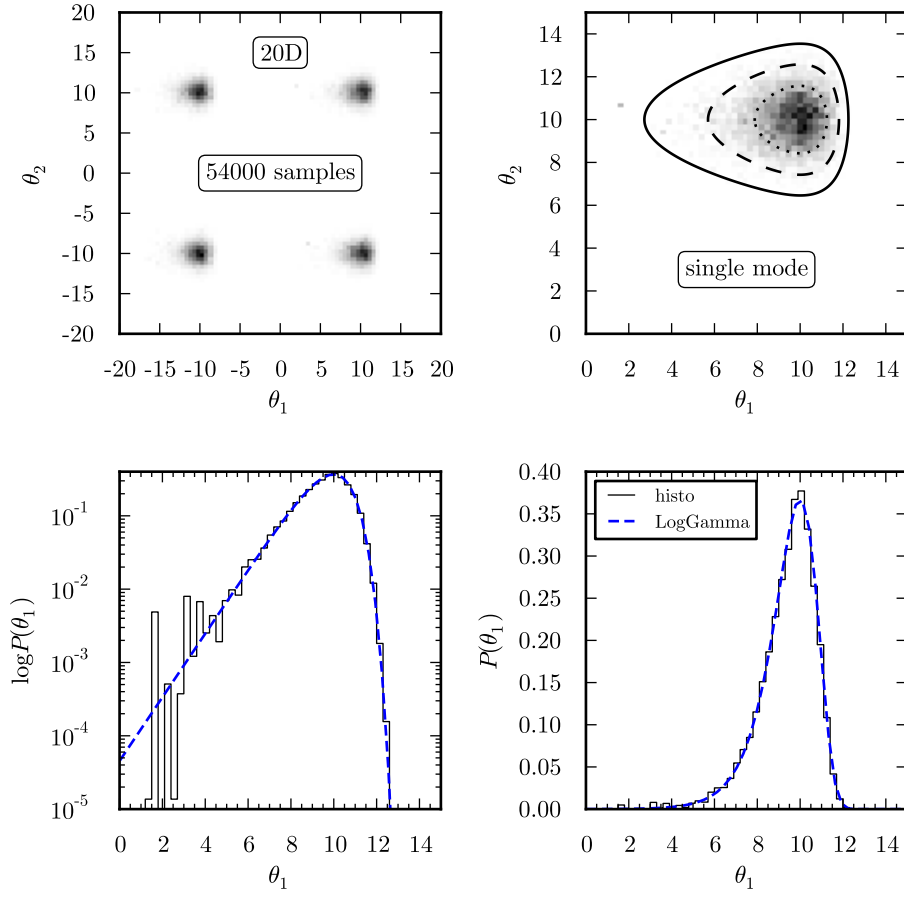


Figure 4: Heavy-tailed mode example in  $d = 20$ . Histogram estimate of  $P(\theta_1, \theta_2)$  for one PMC run with all four modes (top left). Zoom-in on a single mode around  $(\theta_1, \theta_2) = (10, 10)$  with overlaid 1-, 2-, and 3- $\sigma$  contours obtained from integration on a fine grid (top right). The corresponding 1D marginal  $P(\theta_1)$  on the log (bottom left) and linear (bottom right) scale is shown together with the true target  $P(\theta_1) = \text{LogGamma}(\theta_1 | 10, 1, 1)$  overlaid.